Volume Title: The Effect of Pension Plans on Aggregate Saving: Evidence from a Sample Survey

Volume Author/Editor: Philip Cagan

Volume Publisher: NBER

Volume ISBN: 0-87014-409-X

Volume URL: http://www.nber.org/books/caga65-2

Publication Date: 1965

Chapter Title: Saving Ratios of Covered and Not-covered Households Compared

Chapter Author: Philip Cagan

Chapter URL: http://www.nber.org/chapters/c1672

Chapter pages in book: (p. 20 - 54)

3

SAVING RATIOS OF COVERED AND NOT-COVERED HOUSEHOLDS COMPARED

1. Pension and Other Contractual Saving

Some recent studies of saving find that contractual and discretionary saving are substitutes for each other,[1] but give widely differing estimates of the rate of substitution, in part because of small samples that entail large ranges of error. For verification of their results, let us look at the Consumers Union sample and, at the same time, compare the behavior of pension with other contractual saving.

Table 4 gives, for households covered and not covered by group pensions,[2] the average ratios for various kinds of saving and rates of substitution between them. The table makes the exclusions outlined in the preceding chapter, except that the top two rows include households with extreme saving ratios and the bottom two exclude them. In addition, the third row excludes nineteen households with pension

[1] See J. Morgan, "Factors Related to Consumer Saving When It Is Defined as a Net-Worth Concept," in *Contributions of Survey Methods to Economics* (L. Klein, ed.), New York, 1954, p. 135; L. Klein, "The British Propensity to Save," *Journal of the Royal Statistical Society,* Series A, Part 1, 1958, pp. 60–96; and I. Friend and R. Jones, "The Concept of Saving," in *Consumption and Saving* (Friend and Jones, eds.), University of Pennsylvania, 1960, II, 336–359.

On the other hand, the following studies find no clear evidence of substitution. Preliminary results of work by the Survey Research Center show increased saving by people under private pension plans (G. Katona, "Saving for Retirement," *The Outlook on Consumer Behavior* [C. Lininger, ed.], Foundation for Research on Human Behavior, 1964, pp. 103–121). An unpublished Oxford dissertation by N. Liviatan compares saving via British superannuation schemes with other saving and finds little or no substitution. A Swedish study brought to my attention by Thomas Mayer reports a large increase in total saving for households under pension plans (*Hushållens Sparande År 1955* [Household Savings in the Year 1955], Vol. I, Meddelanden Från Konjunkturinstitutet, Stockholm, 1959).

[2] Determined by whether respondents answered any part of question 15 (see Appendix). No attempt was made to correct for households not covered in their present jobs but who reported themselves as covered because they received pension benefits from their previous employment, for example, retired army personnel who took private jobs.

TABLE 4

Average Ratios of Components of Saving to Income and Estimated Rates of Substitution Between Them for Covered and Not-Covered Households

Sample Households[a]	Average Ratios (per cent)				Rate of Substitution		
					Other Contractual Saving with Pension Saving[f]	Discretionary Saving[g] with	
	Discretionary Saving[b]	Pension Saving[c]	Other Contractual Saving[d]	Total[e]		Pension Saving	Other Contractual Saving
Excluding those with gains or losses over $1,000							
Covered[h]	2.8	3.2	11.6	17.6	$.25^s$ (±.05)	$.10^s$ (±.03)	$-.23^s$ (±.01)
Not covered	2.3	—	9.1	11.4	—	—	$-.28^s$ (±.03)
Excluding those with gains or losses over $1,000 and extreme saving ratios[i]							
Covered[h]	2.8	2.8	5.9	11.5	$.28^s$ (±.06)	$.21^s$ (±.06)	$-.73^s$ (±.02)
Not covered	2.1	—	5.7	7.7	—	—	$-.77^s$ (±.03)

Note: Figures in parentheses give the range of error at the .05 level of significance; s means coefficient is significantly different from zero at that level.

[a] Self- and not-employed excluded. Number of households included: 1st row, 8,414; 2nd row, 3,099; 3rd row, 8,008; 4th row, 2,911.

[b] Reported increase in cash and securities minus increase in nonmortgage debt.

[c] Reported increase in equity in pension fund.

[d] Reported increase in equity in life insurance and annuities (cash surrender value) and real estate (but not in pension fund).

[e] May not exactly equal the sum of figures shown because of rounding.

[f] Estimated regression coefficient when other contractual saving ratio is regressed on pension-saving ratio.

[g] Estimated partial regression coefficients when discretionary saving ratio is regressed on pension saving ratio and other contractual saving ratio.

[h] Here and subsequently in this study coverage under a group pension plan other than OASDI was determined by whether respondents answered any part of question 15 (see Appendix).

[i] Excludes households with pension saving ratios of 50 per cent or more and with other saving ratios of 50 per cent or more in absolute value.

saving ratios over 49 per cent who probably reported erroneously. The exclusion of extreme ratios affects the size of the figures but not the direction of effects in which we are interested. As argued in Chapter 2, this exclusion seems appropriate, and we may assume that the bottom two rows give the more reliable estimates. For these, the last column shows that discretionary saving is less by 73 and 77 cents for every dollar of other contractual saving, which is in line with other studies. The covered and not-covered households do not have materially different rates. The next two columns show, however, that pension contributions have quite different effects. For every dollar of such contributions there is 21 cents *more* of discretionary and 28 cents *more* of other contractual saving; these rates differ from zero at the .05 level of statistical significance.

A positive rate is confirmed by the higher average ratios for covered than for not-covered households: the former save more in total and do so by more than the amount of their pension contributions. There is, to be sure, a minor inconsistency between the rates and the average ratios. By the regression equation, the rates for covered households in the third row imply that pension contributions of 2.8 per cent of income increase other contractual saving by ($2.8 \times .28$) 0.8 percentage points, whereas the actual increase over the not-covered households in the fourth row is only 0.2 points. On the other hand, the rates and averages for covered households imply an increase in discretionary saving of ($2.8 \times .21 + 2.8 \times [1 - .73] \times .28$) 0.8 percentage points, and the actual increase over the not-covered households is 0.7 points. Some discrepancy is not surprising; there are bound to be random variations in the data, particularly in the average ratios for not-covered households, which number only about a quarter of our sample. Another reason for inconsistency may be that the effect of pension on other saving is not linear, contrary to the assumption implicit in these regressions. Evidence of nonlinearity is presented in Chapter 4.

Pension saving reported by households is likely to be inaccurate, since many respondents may be uncertain about the amount of their own contribution and probably have only the haziest notion of their employer's.[3] The average pension contribution of 2.8 per cent of income seems far too low compared with our separate estimate of it

[3] Most industrial plans have fixed benefits, but do not specify the amount of contributions by employers. The plan for many college teachers is atypical in having

(see Chapter 4, Section 2) and other information. Over two-thirds of the respondents either made no estimate of the change in pension equity or reported that their estimate included no part of the employer's contribution. Of those that made no estimate, the majority had noncontributory plans. Nonetheless, these data give the respondents' understanding of what their pension saving is. From this point of view, the results indicate that coverage does not lead households to substitute their pension contributions for other forms of saving, but seems actually to induce a slight increase in other saving. If so, aggregate personal saving is increased by the amount of growth in pension funds and apparently also by a small amount of increased saving by covered households in other forms. Such results suggest one reason why other studies obtained wide-ranging estimates of the rate of substitution between discretionary and contractual saving: they have included in the latter a large and variable item—pension contributions—that has a quite different rate.

An immediate objection to these results is their neglect of social security coverage (OASDI). Eligible households will take anticipated social security benefits into account in their retirement plans. It is possible that a greater fraction of the no-pension group is under OASDI, taxes for which are excluded from these saving figures. This neglect is relevant to our sample: social security covers about 80 per cent of the pension group and about 94 per cent of the no-pension group (in part because the former includes a large group of government workers who are not covered by social security). The pension group may save more in other forms only because a larger fraction does not have social security, and this possibility can be checked by excluding households that are not so covered. The relevant figures are presented in Table 5 for the reduced sample of households covered by a pension in the third row of Table 4. The corresponding figures for the no-pension group are not shown because that group is almost entirely covered by social security.

The rates of substitution in Table 5 follow the pattern in the pre-

---

stated contributions by both employee and employer, up to half of which may be paid into a variable annuity, CREF, the rest to a fixed-dollar annuity, TIAA. Pension payments from CREF are determined by the investment experience of its common stocks. Knowledge of the employer's contribution is discussed further in Chapter 5.

TABLE 5

Average Ratios of Components of Saving to Income and Estimated Rates of Substitution Between Them for Households with Pension Plans[a] by Social Security Coverage[b]

Households with Pension Plans	Average Ratios (per cent)				Rate of Substitution[c]		
					Other Contractual Saving with Pension Saving	Discretionary Saving with	
	Discretionary Saving	Pension Saving	Other Contractual Saving	Total		Pension Saving	Other Contractual Saving
With OASDI	3.1	2.6	5.9	11.6	.30^{s}(±.07)	.26^{s}(±.08)	−.71^{s}(±.03)
Without OASDI	1.7	3.7	5.9	11.3	.22^{s}(±.13)	.17^{s}(±.10)	−.82^{s}(±.04)

Note: Figures in parentheses give the range of error at the .05 level of significance; s means coefficient is significantly different from zero at that level.

[a] Same exclusions as Table 4, third row, as well as 169 households with pension plans for whom social security coverage could not be determined. Number of households included: 1st row, 6,297; 2nd row, 1,542.

[b] Based on question 2 (see Appendix).

[c] See Table 4.

ceding table: negative for the effect of other contractual saving on discretionary saving and significantly positive for the effect of pension saving on discretionary and other contractual saving. The households with OASDI even show the positive effect of pension saving more strongly, though the differences between the estimated rates for the two groups may well reflect random variability.

The discretionary saving ratio and the rates of substitution for pension saving are higher for the households with OASDI than without, which seems to imply that social security coverage leads to greater saving in other forms, just as we find for pension saving. Such an effect is often attributed to social security, but it cannot be inferred from our data. Most of the households in our sample without OASDI work for governmental agencies (including teachers in public education). All that the comparison in the first column of Table 5 shows is that government workers save less in discretionary forms than others do, which may be due to their lack of social security or to some other common characteristic. (One is suggested in Chapter 4, Section 1.) Also, in government the workers make about one half of the total contribution to their plans, which is uncommonly high and explains why the pension saving ratio is higher in the second than in the first row of Table 5 if, as seems likely, employers' contributions were typically not reported.

This occupational distortion in the sample of the effect of social security coverage does not apply to our interpretation of pension saving, and we may conclude that this kind of saving is associated with a slight increase in other saving and that this is not due to differences in social security coverage. The next step is to check the data for other possible distortions and, if none is found, to see what explanation can be adduced for this association.

2. *Adjustment for Age and Income*

Pension saving as reported is, as has been noted, suspect because of its uncertain meaning and accuracy. For the remainder of this chapter, let us disregard the reported pension contribution in question 14 altogether and look only at total *other* saving, which excludes it. The preceding results can be checked by comparing the other saving ratios of households divided into covered and not-covered groups. Table 4 has shown that covered households save more than not-covered households *in addition to* pension contributions—0.7 percentage points more in discretionary saving[4] and 0.2 more in contractual saving, or 0.9 more in total, according to the averages in the bottom two rows of the table. (The increase in contractual saving indicates that pension saving does not substitute for saving via life insurance, which seems a priori to be the closest substitute.) The range of error of this total difference between the two groups is ±0.6 at the .05 level, so it is statistically significant despite its small amount. We need, however, to check for possible biases.

One possibility is that the heads of the covered households are likely to be older and better-paid workers. This, combined with the tendency of saving ratios to rise with age, might explain the greater saving of the covered group. In addition, though higher-income groups do not necessarily save proportionately more, they might. One way to adjust for age and income is to compare households in each age (according to head) and income class separately. This procedure leads to comparisons between small subclasses of the sample, however, and so increases the errors of estimation. The same purpose can be achieved with a three-way analysis of variance, which is equivalent to estimating the average difference in saving ratios between the various age and income classes, adjusting the ratios for these differences, and then finding the average difference in these adjusted ratios between the covered and not-covered groups.[5]

[4] Cash saving (that is, "checking and saving accounts, government bonds") is responsible for all the difference in discretionary saving. The breakdown for the bottom two rows of Table 4 is as follows:

	Saving Ratios (per cent)		
	Cash	Securities	Total Discretionary
Covered	2.08	0.69	2.77
Not covered	1.28	0.78	2.06
Difference of covered over not-covered group	0.80	−0.09	0.71

[5] The method is equivalent to fitting the following regression equation to the data by least squares:

$$\begin{aligned} S/Y = k &+ (a_1 - a_2)A_1 + (a_3 - a_2)A_3 + \cdots + (a_6 - a_2)A_6 \\ &+ (y_1 - y_3)Y_1 + (y_2 - y_3)Y_2 + (y_4 - y_3)Y_4 + \cdots + (y_7 - y_3)Y_7 \\ &+ (c_1 - c_2)C \end{aligned}$$

where S/Y is the ratio of total other saving to income, and A_i is unity if the house-

The analysis of variance simplifies interpretation of the results by giving one estimate of the effect of pension coverage based on the full sample. It has the minor drawback of assuming no interaction effects between age and income (that is, that the effect of age on the saving ratio does not depend on the income class nor the effect of income on the age class). Such interactions are not likely to be important, at least compared with those between pension coverage and the age or income class (discussed later). This method has the advantage over standard multiple regressions of not assuming any particular functional form for the effect of age and income.

Table 6 gives the results of this three-way analysis of variance for the full sample and the middle age and income classes separately. Corrections for any two age or any two income classes show the mean difference between the two classes standardized for the effects of the other regression variables. As they stand, the figures give standardized mean differences between each class and the arbitrarily selected standard class (25–49 for age [6] and $5,000–$7,500 for income). That is, the estimated saving ratio for not-covered households age 25–49 and income $5,000–$7,500 is the constant term (given at end of footnote 5). The estimated ratio for any other class is found by adding the appropriate correction factors. Most of the factors are statistically significant.

The effect of pension coverage standardized for age and income is

---

hold is in the ith age class and zero otherwise, Y_i is unity if the household is in the ith income class and zero otherwise, C is unity if the household is covered by a pension fund and zero otherwise, and the small letters are regression coefficients with the following interpretation:

$a_i - a_2$ is the difference in average saving ratios between the ith and 2nd age classes (the latter being arbitrarily selected as the standard age class) standardized for income and pension coverage, $y_i - y_3$ is the difference in average saving ratios between the ith and 3rd income classes (the latter being arbitrarily selected as the standard income class) standardized for age and pension coverage, $c_1 - c_2$ is the difference in average saving ratios between the covered (c_1) and not-covered (c_2) households standardized for age and income, and k is the mean ratio for not-covered households in the 2nd age class and 3rd income class. In the regressions of Table 6, k was 6.7 ($\pm$0.7) for all classes and 6.8 ($\pm$0.7) for the middle classes.

[6] This class was made as broad as it is on the basis of preliminary computations, which revealed little difference between subclasses within the 25–49 boundary. However, the negative correction for the 50–54 age class, which on a priori grounds seems to have the wrong sign, may reflect the upward pull on the ratio for the 25–49 class of the younger age groups in that class, which in this sample generally seem to have comparatively high ratios (for what reason is not known).

shown in the bottom row. The pension effect is smaller here than in Table 4, in which age and income effects were ignored, and according to the range of error could be as small as zero or as large as 1.2. The preceding conclusions are therefore confirmed in that covered households make no net reductions in other forms of saving and in all likelihood make a net increase, though it may be less on the average than one percentage point.

TABLE 6

Corrections for Ratio of Total Other[a] Saving to Income to Standardize[b] for Age, Income, and Pension Coverage

Standardized Classes	All Age and Income Classes[c]	Middle Age and Income Classes[d]
Age Class		
1. Under 25	+4.0[s] (±2.5)	--
2. 25 - 49	0	0
3. 50 - 54	-1.1[s] (±1.0)	-1.1 (±1.1)
4. 55 - 59	+0.6 (±1.2)	+0.6 (±1.3)
5. 60 - 64	+2.3[s] (±1.7)	--
6. 65 and over[e]	+3.1[s] (±2.6)	--
Income Class		
1. Less than $4,000	-2.5[s] (±1.9)	--
2. $4,000 - 4,999	-2.6[s] (±1.3)	-2.9[s] (±1.4)
3. $5,000 - 7,499	0	0
4. $7,500 - 9,999	+1.2[s] (±0.7)	+1.2[s] (±0.7)
5. $10,000 - 14,999	+2.6[s] (±0.7)	+2.6[s] (±0.7)
6. $15,000 - 24,999	+4.8[s] (±1.1)	--
7. $25,000 and over	+5.6[s] (±2.5)	--
Covered Class	+0.6 (±0.6)	+0.4 (±0.7)

Note: Figures in parentheses give the range of error at .05 level of significance; s means correction is significantly different from zero at that level.

[a] That is, excluding contributions to group pension plans and social security taxes.

[b] Estimated by the method of fitting constants, which is equivalent to fitting a regression equation to the data by least squares. See footnote 5, above, and G. Snedecor, *Statistical Methods*, Ames, Iowa, 1946, 4th ed., pp. 296–299.

[c] Reduced sample of 10,938 (of which 8,027 have pension plans) after exclusions described in Chapter 2; same as Table 4, bottom two rows together, except that the nineteen households with extreme pension saving ratios are included here.

[d] Reduced sample of col. 1 with ages under 25 and over 59, and incomes less than $4,000 and over 14,999, omitted.

[e] Employed only.

Covered households tend to have higher incomes, and this remains true within income classes. Any effect produced by income levels might still distort these figures. Yet, while there may be slight distortion, Appendix Table B, which gives average income and saving for each income class in Table 6, shows that the higher saving of covered households within most classes is far more than can possibly be explained by their higher income. (Possible differences in disposable income are discussed in Section 5 of this chapter.)

The extreme age and income classes, which have ratios considerably higher than the middle classes (except for the lowest income class), do not greatly affect the over-all results. The analysis for the middle classes alone, shown in the second column of Table 6, gives a +0.4 pension effect, which is a little lower than for the full sample.

The extreme classes are relatively small in number and for that reason can be expected to have little effect on estimates for the full sample. The extreme classes nevertheless save quite differently than do the middle classes. To determine the difference, it is necessary to analyze them separately. For this analysis the lowest age and income classes are excluded because their saving is very likely subject to special influences and, in any event, they comprise few households. With their exclusion in addition to the middle classes, there remain the higher age or income classes excluded from the regression for the second column of Table 6. Since this remainder covers a small part of the sample, age and income correction factors estimated for it separately would be unreliable. We may apply the factors estimated for the full sample and estimate from this remainder just the standardized difference between the covered and not-covered groups. Thus computed, the estimate of this difference is 0.6 (±2.0). It is identical to the estimate for the full sample, suggesting that the pension effect does not vary with age or income, though the range of error is large because of the comparatively small numbers on which the estimate is based. In statistical terms, the equality of these two estimates of the pension effect suggests that the interaction effects between coverage and age or income are zero, or at least that this is so as between coverage and extreme values of age or income. Tabulations (not presented) indicate that the zero interaction does *not* result from large effects by age and income that happen to offset each other.

The absence of an age effect is surprising and important. One might have expected that older workers would be sensitive to the circumstances of their retirement and would respond strongly to changes in the preparations made for it. On this expectation, today's older workers would display a sophistication regarding pension plans that tomorrow's households could be expected to acquire at early ages, most having by then the benefit of the experience of fathers and uncles who had learned how the plans work at first hand. Insofar as this sophistication does not increase the effects of coverage, as the absence of interaction suggests, these effects on an individual basis will not change over time, and their strength on an aggregate basis will continue to rise as pension funds grow, at least in the near future.

One difficulty with these tabulations is that they take no account of households with more than one income earner. With an extra earner (typically the wife), the household is more likely to have at least one member covered by a pension plan. A positive association between the saving ratio and coverage would then result if a disproportionately large fraction of the income of extra earners was added to saving (in anticipation perhaps of large future purchases, or to pay off debts).[7] We can easily check on this possibility by segregating our sample by the number of income earners per household. The results are shown in Appendix Table C. Most households have one or two income earners; in this table the few having three or more have been grouped with those having two. The covered households still have significantly higher nonpension saving both among households with one earner and among those with two or more earners.

Another drawback of the adjustment for income level made in the foregoing tabulations is the use of current actual income instead of each household's long-run expected or permanent income level. The two can differ and thereby introduce error into the results because current income may be temporarily high or low owing to unusual

[7] I am indebted to Margaret Reid and Milton Friedman for this point. See W. J. Bilkey, V. G. Massaro, and J. P. Meehan, Jr., "The Structural Effects on Consumer Disbursements of Wives Working," *Review of Economics and Statistics,* May 1962, pp. 221–224, which suggests that working wives produce higher household saving ratios, at least in the early and middle years. On the other hand, our Consumers Union sample shows multiple-earner households to have slightly lower saving ratios (see Appendix Table C).

circumstances, or because current income has yet to reach (or will soon fall from) the long-run level the household expects during its middle years or when present training and apprenticeships are completed. In line with Milton Friedman's permanent income hypothesis,[8] households may be expected to save a large part or all of positive deviations of current from permanent income and to dissave when confronted with negative deviations. If so, the ratio of saving to permanent income gives a better measure of households' normal saving habits, and the use of current income as a substitute may distort the comparison.

The distortion would be serious here only if the covered and not-covered households differed in the prevalence of these deviations from permanent income. Since this might occur by accident, we should test for it. Although we have no estimate of permanent income levels, we can adjust for one of the causes of such deviations—transitory changes in income level. This is done in Appendix D by excluding all households that reported a temporary rise or fall in income. The results are essentially the same as in Table 4, which did not make these exclusions, suggesting that the drawbacks of current income as a proxy for permanent income have not affected the comparison.

Excluding households with transitory changes in income, of course, takes no account of unfortuitous deviations from the long-run expected level. Such deviations would occur particularly among young persons with advanced education, who must put off the full rewards of professional standing during long periods of training and apprenticeship. This is typically true of doctors, most of whom we exclude in omitting the self-employed. But it is equally true of salaried engineers and other college graduates, whom our sample includes. These people know that in the early years their current income will be far short of the levels to be attained later, but they would not report the early deficiency as a "transitory fall"; a steady rise in their income level is fully expected and cannot be viewed as transitory up or down. If such groups were represented unequally among the covered and not-covered households, the preceding results would be distorted. Grouping by education and occupation will eliminate most of this bias.

[8] *A Theory of the Consumption Function,* Princeton University Press for NBER, 1957.

3. Adjustment for Education and Occupation

Our sample has a high proportion of teachers and government workers—special occupational groups that may have atypical saving habits. The sample is also heavily weighted with college graduates, who appear to have higher saving ratios than less-educated workers, as suggested by this study and others. If the covered and not-covered groups were represented differently among occupations and educational levels, the comparison would be distorted. Also, a breakdown by occupation, as already suggested, brings together households likely to experience similar patterns of lifetime income and so eliminates some of the bias from using actual rather than long-run expected income levels.

The choice of specific occupations for classification was largely dictated by the special character of the sample. The first step was to segregate salaried workers from wage earners, because, even aside from differences in relative income level, the two groups differ in the stability of their income over time and perhaps also in general attitudes toward saving; this is the traditional distinction between white- and blue-collar workers. The next step was to segregate the salaried employees of educational (including religious and social service) [9] institutions and of the government, which constitute a substantial part of the sample and conceivably have saving habits not representative of the average American household. After these groupings, the remainder of the sample consisted of one large class covering manufacturing, trade, transportation, and finance and several small related classes. They were left as an all-inclusive "business" group. The latter includes salaried workers in all types of business firms, employees on commission, salaried workers in agriculture, and a small number of respondents who work for labor unions. A small number of salaried workers in manual jobs were omitted for want of clear evidence whether they were appropriately designated white- or blue-collar workers. All wage earners, on the other hand, appear to be production workers or the like in business firms. This group was relatively small, in part because it

[9] Clergy and social service were grouped together as a category on the questionnaire, but had so few households here (see Appendix Table E) that they were included for analysis with teachers, which seemed to be the most similar occupation for our purposes.

is unrepresentative of Consumers Union subscribers and in part perhaps because many people tend, if in doubt about the question's meaning, to upgrade themselves in the occupational hierarchy.

Occupation is closely associated with educational level, and differences in saving behavior among occupational groups may partly reflect differences in education, and vice versa. It is desirable to classify by both characteristics. Question 3 elicited the educational attainment of the head of each household in terms of four levels: high school attended or completed, college attended, four-year college completed, and graduate school attended or completed. The last two groups were combined, since preliminary tabulations showed that they had nearly the same average saving ratios.

The cross classification is presented in Table 7, in which some cells involved small numbers and were omitted. The table shows higher saving by the covered groups, though most of the differences are not statistically significant. The preponderance of positive differences is significant, however, as shown by the sum of the differences at the foot of the table. Hence this classification of the data reaffirms the previous findings.

Comparing the various groups, it is seen that, aside from coverage, the higher educational groups save more; that among occupational groups, business and government have larger savers than the education, clergy, social service group; and that in this company the wage earners are the smallest savers. Indeed, the very low saving ratio for the not-covered wage earners produces the two largest differences in the table between the covered and not-covered housholds. Much of the variation in the saving ratio between educational and occupational groups is probably not statistically significant, and so too much should not be made of it. Nevertheless, it does conform to the results one finds in classifications by income: this and other cross-sectional evidence suggests that higher saving ratios go with higher income, more education, and a professional or salaried occupational status; however, it is difficult to disentangle cause and effect among the three variables.[10] The

[10] We may note in passing that the effects of these variables imply a secular rise in the aggregate saving ratio, unless a household's relative rather than its absolute standing in the population is important; whereas, in fact, the secular trend appears to be flat. The reason might be that these cross-sectional differences are too small

TABLE 7

Average Ratios of Total Other[a] Saving to Income of Covered and Not-Covered Households by Educational and Occupational Group[b]

Educational and Occupational Group[c]	Average Saving Ratio (per cent) Covered	Not Covered	Difference of Covered over Not-Covered Group[d]
High-School Graduate or Less			
Salaried employee of			
Business	8.0	8.6	-0.7 (±2.2)
Wage earner	7.2	4.4	2.8[s] (±2.2)
Some College			
Salaried employee of			
Business	8.8	8.3	0.5 (±1.6)
Government	7.9	6.8	1.2 (±5.1)
Wage earner	7.0	4.3	2.7 (±3.0)
College Graduate or More			
Salaried employee of			
Business	10.5	9.0	1.6[s] (±1.1)
Government	8.6	9.7	-1.1 (±3.0)
Education, clergy, social service	7.7	6.0	1.7 (±1.7)
Sum of differences for all groups above	—	—	8.6[s] (±7.9)

Note: Figures in parentheses give the range of error at the .05 level of significance; s means significantly different from zero at that level.

a That is, excluding contributions to group pension plans and social security taxes.

b Same exclusions as for Table 6 (see note c of that table) as well as households in groups not listed above. Numbers of households in covered and not-covered groups respectively: 1st row, 476 and 264; 2nd row, 447 and 259; 3rd row, 890 and 497; 4th row, 373 and 31; 5th row, 222 and 140; 6th row, 2,314 and 1,008; 7th row, 940 and 85; 8th row, 1,334 and 319.

c Education, clergy, and social service are not shown in the first two educational groups, government is not shown in the first group, and wage earners are not shown in the last group because the number of households was too small to make the comparison worthwhile.

d May not exactly equal the difference between the ratios shown because of rounding.

saving pattern among cells in the table therefore appears to be representative of the nation's population, and the observed difference between the covered and not-covered groups seems to have general applicability.

To check for the effect of differences in income, let us examine the average level of income and saving for each educational and occupational group (Appendix Table F). Covered households tend to have higher incomes, as was to be expected. But the differences are too small to explain the differences in saving: in most cells the higher average saving of the covered over the not-covered group is much more than can be explained by the slight difference in income.

The figures in Table 7 are consistent with the possibility (alluded to in Chapter 2) that the economic recession during the survey period might have raised the saving ratios of the covered more than the not-covered households and that the differential rise would have been sharpest for wage earners. This explanation is based on the greater insecurity of certain jobs during recessions and so is not likely to account for the positive differences for the government and education, clergy, and social service groups, but it might apply to the business employees and wage earners, and so account for the observed absence of reductions in the other saving of covered households. We may consider this possibility better from a breakdown of households by the cyclical sensitivity of the industry in which they were employed. Appendix Table E presents a rough classification of this kind and shows that covered households still save more by an amount consistent with the other comparisons. Distortions produced by the particular time period in which the survey was conducted, therefore, appear to have been negligible.

4. Ratio of Wealth to Income

One difficulty with these results is that they may reflect merely a tendency for households with high propensities to save to gravitate toward jobs offering pension plans. Since they segregate covered and not-

---

to affect the secular trend or that they change their direction of influence at high levels so far achieved by only part of the population. The proper interpretation of these differences is not relevant to this study and need not be pursued here.

covered households, the preceding tables may not, as is implicitly assumed, compare households having equal propensities to save; they may thus be misleading in regard to the effect of coverage. Although few people are likely to select their job primarily on the basis of its pension benefits, since there are other aspects of employment many times more important, one cannot dismiss this possibility altogether. If a small fraction of workers who were above-average savers preferred jobs with pension plans, they could produce the results observed. Workers who seek employment with large corporations offering pension plans may do so to satisfy a desire for security and so may tend to save more; this is consistent with the data, but invalidates the previous interpretation. On the other hand, such workers may achieve the sought-after security when employed by large corporations, in part because pensions are offered, and so may tend to save less; this seems to conflict with our data, however. Or there may be other differences between the covered and not-covered groups that affect their saving and distort the comparison.

One slight piece of evidence already presented, however, denies a job-selection process at work in the sample. In Section 2, above, it appeared that coverage and age do not interact; that is, the difference in average saving ratios between covered and not-covered households was the same for the older age classes as for all. This finding is relevant because the selection process described above should affect the younger households most. Pension plans have spread rapidly since World War II, and new workers entering the labor force over the past decade had a good opportunity to take jobs with pension plans if they wished. An older worker who wanted a pension plan but whose employer had not started one could, of course, have changed jobs. However, since mobility is much lower among older than younger workers, the latter should show the higher association between coverage and propensity to save. Not finding such differences among age classes suggests that this selection process is weak or nonexistent.

Further evidence against the job-selection explanation was provided by Table 4. There a positive relation was found between discretionary saving and the reported amount of contribution to pensions within the covered group by itself. Such a relation cannot be explained by the job-selection hypothesis. Yet, though this relation implies an effect that varies with the amount of contribution, the importance of coverage itself cannot be inferred directly. Moreover, if one placed great importance on differing individual propensities to save, he might contend that this observed relation is spurious: it could be that covered households with a high propensity to save were for that reason more aware of their own and their employers' contributions and reported them in full. Households having weaker saving propensities may have forgotten to report or underreported their own pension contributions and omitted their employers'. In that way the observed relation could be produced with no implication that a higher contribution increases other saving. This possibility appears to be remote; still, the mere fact that it is a possibility appeals for other evidence to clarify the question.

Another test of whether there are extraneous differences between the covered and not-covered groups in the sample can be made by using data on total wealth, which the second survey elicited. Wealth reflects past saving; since we also know the number of years each household has been covered by a plan (question 15C), and hence from the age of its head the number of years not covered, we can by regression methods estimate for a group of previously covered households their average saving for years covered and for years not covered.[11] Since these estimates are based on the group covered at the time of the survey, a comparison with the not-covered group is avoided. In addition, estimates of saving based on wealth pertain to many years and thus supplement from independent evidence the previous estimates based on one particular year.

Equity in pension plans (as well as retirement benefits under social security) is excluded from total wealth to give total other wealth, which is comparable to total other saving used in the foregoing sections. It is (see question 13) the addition of:

Checking accounts
Saving accounts and government bonds
Common stock and mutual funds (current market value)
Other marketable securities

[11] A problem here and elsewhere that we have neglected is that, in covered households with two or more income earners, the head may not be the one with a pension plan. Because of the way the questionnaire was worded, there is little we can do about this problem other than hope it is not serious.

Market value of home
Equity in annuities and life insurance (cash surrender value)
Other assets (specify) [12]

minus

Mortgage on home
Other personal debt (instalment, etc.)

Respondents checked one of the following boxes for each item: don't have, under $500, $500–1,000, $1,000–2,000, $2,000–5,000, $5,000–10,000, $10,000–20,000, $20,000–40,000, and, if over $40,000, wrote in the amount. The midpoints of the checked boxes or the amount written in were added (or subtracted for debt) to get the total.

Any figures that purport to measure wealth undoubtedly contain gross errors. Although the sample comprised respondents above average in willingness to be careful and accurate, the question on wealth surely taxed the patience and understanding of even the most conscientious of them. The wording of the items was cryptic and ambiguous (from the necessity to avoid long and complicated questions). On the brighter side, one item particularly hard to answer—equity in pension plans—was not used. Errors of reporting the other items enlarge the variance of total wealth, but do not seem to bias the estimates we want to derive. Using midpoints probably biases the total downward, but this, too, seems unimportant for our purposes.

The most serious difficulty is that wealth is acquired in two ways. In addition to saving out of income, it reflects gifts, capital gains, and inheritances. We can still measure past saving with reasonable accuracy if the fraction of wealth acquired the second way is relatively low and distributed at random among households. Because of these nonsaving additions to wealth, estimates of saving derived from wealth will be inexact, but estimates from large samples should still be relevant and useful.

To make comparisons with the saving ratios used previously, the ratio of wealth to income was computed. A wealth-income ratio of 1.0, for example, implies average saving for ten working years of 10 per

[12] Real estate, business assets, and personal items (the latter omitted if over $5,000). This category is minor and, though excluded by our preceding concept of saving, should not greatly affect the comparability of the wealth and saving figures.

cent, or for twenty working years of 5 per cent, of the latest year's income. As is discussed further later, this may understate an average of past saving ratios if income, as it usually does, rises with age. (Of course, wealth contributes to current income through interest earnings, but the contribution is much smaller than wages or salary for most employed households.) Average wealth ratios are presented in Table 8 by age group and by years covered under a plan. To avoid the distortion of extreme ratios, seventy-two households with ratios of 10 or greater were arbitrarily excluded.

TABLE 8

Average Ratios of Total Other[a] Wealth to Income by Age Group and Period Under a Pension Plan[b]

Age	Not Covered	Years Covered Under a Pension Plan: Less Than 5	5 to 9	10 to 14	15 and Over
Under 25	1.3	0.8	0.5	c	3.8[d]
25-29	1.2	1.0	1.1	1.3	4.0[d]
30-34	1.4	1.2	1.3	1.6	1.5
35-39	1.8	1.6	1.6	1.9	1.8
40-44	2.2	2.0	1.9	2.1	2.1
45-49	2.4	2.2	2.5	2.4	2.5
50-54	2.7	2.3	3.0	2.8	2.9
55-59	2.8	2.8	2.9	2.9	3.1
60-64	3.4	3.0	3.4	3.1	3.4
65 and over[e]	4.3	3.9	3.6	3.9	3.8
All age groups	1.8	1.5	1.7	2.2	2.6

[a] That is, excluding equity in group pension plans.

[b] Same exclusions as for Table 6 (see note c to that table) as well as questionnaires deficient in information on wealth and years covered or having wealth-income ratios of 10.0 or greater. Number of households in bottom row: 1st col., 2,824; 2nd col., 2,932; 3rd col., 2,016; 4th col., 1,387; 5th col., 1,472.

[c] No households in cell.

[d] Apparently involves error in reporting length of coverage. The number of households is too small to affect the results.

[e] Employed only.

In comparing the rows, one first notices that the not-covered group has higher wealth ratios on the average than the households covered less than ten years. The difference for all age groups is statistically significant for the not-covered and covered-less-than-five-years groups at the .05 level, but not the not-covered and covered-less-than-ten-years groups. One might infer that the higher wealth ratio for the not-covered than for the recently covered households reflects greater past saving by the former, but this is not consistent with the rest of the table. Among the covered groups, a higher wealth ratio is generally associated with longer coverage, thus suggesting greater saving when covered than when not. A possible explanation of the higher average ratio for the not-covered than for the recently covered households is that the former inherited more, though why they did is not clear.[13]

The not-covered group aside, we can estimate the effect of coverage from the covered group alone. While the covered households seem to differ from the not-covered group in the amount of inherited wealth, they are less likely to differ in this way greatly or systematically among themselves. All we need to assume is that the age at which households covered at the time of the survey acquired their plans in the past was not related to their propensity to save; that is, that households with higher propensities to save did not systematically become covered at younger ages. This assumption can be tested by comparing the current saving ratios of households covered for different lengths of time; the comparison is made in Chapter 4, Section 3, and shows no tendency for longer periods of coverage to be associated with higher saving ratios.[14] On this basis Table 8 shows that, within age groups, the wealth ratio rises on the whole with length of coverage, which implies that households saved more when they were covered than when not.

[13] One reason could be that income in preceding years rose faster for the covered group, which would make past saving lower and the ratio of accumulated wealth to present income lower. This was likely to have been true for government workers and teachers, most of whom were under pension plans. I am indebted to Robert J. Landry for this point.

Another reason might be that the not-covered group includes many people who have great wealth invested in their own business, but who draw regular salaries and were not classified as self-employed, such as, perhaps, owners of automobile distributorships or other independent businessmen who said most of their income came from salary. None of them is likely to be covered by group pension plans. A small number of such people who had wealth-income ratios less than 10 and so were not excluded might raise the average wealth ratio substantially, but the average saving ratio very little.

[14] This comparison does not hold age constant, but it is not invalid for that reason. In so far as age and length of coverage are correlated (positively), the saving ratio would tend to rise with length of coverage owing to the effect of age, which therefore cannot invalidate the inference drawn in the text from these data.

How much more they saved can be estimated by regression methods.[15] The wealth ratio of each household at the time of the survey reflects the initial level in the first year of full-time employment (assume age 22 for all households),[16] and the saving ratios, when covered and when not, increased by the number of years applicable to each. Let W/Y be the wealth ratio at the time of the survey and $(W/Y)_{22}$ at age 22. Let S_c and S_{nc} be the average saving ratios when the household was covered and when not, and P_c and P_{nc} the respective number of years of each. $P_c + P_{nc}$ will be the age of the head of the household minus 22. The regression equation is

$$\frac{W}{Y} = \left(\frac{W}{Y}\right)_{22} + S_c P_c + S_{nc} P_{nc}.$$

The same equation can be fitted separately to the not-covered households, for all of whom P_c, of course, is zero.

W/Y is the dependent variable, representing the data underlying Table 8. $(W/Y)_{22}$, S_c, and S_{nc} are regression coefficients, estimated by fitting the equation to the data. P_c and P_{nc} are the independent variables, approximated by the midpoints of the age and years-of-coverage intervals. If a household was in the 45–49 age group, the head of the household was taken to be 47 and to have worked 47–22, or 25 years. If he had been covered 5–9 years, P_c was taken to be 7 years; and P_{nc} to be 25–7, or 18 years.[17] For the open-end age group, the midpoint was arbitrarily assumed to be 67; and for the open-end years-covered group, 17 years. To check the possibility that this last midpoint is inaccurate and distorted the results, computations (not shown) excluded the households in that interval; these gave virtually the same estimates.

Estimates of the regression coefficients by least squares and their range of error are presented in Table 9. The saving ratio for the group covered at the time of the survey is estimated to have increased by 2.1 percentage points when the household acquired a pension plan, statis-

[15] The method followed here was described in Cagan, "The Use of Wealth to Compare Households' Average Saving," *Journal of the American Statistical Association*, September 1964, p. 737.

[16] Estimates (Table 9) of saving ratios from the regression equation are not affected by the particular starting age chosen; only the constant term of the equation, the initial wealth ratio, is affected.

[17] P_c and P_{nc} are of course negatively correlated in the sample. The effect of this is to increase the standard error of the regression coefficients, but not to bias them, except for the special reasons noted later.

TABLE 9

Regression Estimates of Saving Ratios for Covered and Not-Covered Years from Wealth-Income Ratios of Covered and Not-Covered Households at Time of Survey [a]

Pension Status at Time of Survey	Ratio of Other[b] Wealth to Income at Age 22	Ratio of Other[b] Saving to Income (Per Cent): Covered Period (S_c)	Not-Covered Period (S_{nc})	Difference of Covered over Not-Covered Period[c] (S_c-S_{nc})
Covered	.61 (±.08)	8.2	6.2 (±0.4)	2.1 (±0.7)
Not covered	.85 (±.11)	--	6.3 (±0.6)	--

Note: Figures in parentheses give range of error at .05 level of significance (error for covered period not computed).

a Least-squares estimates using equation in accompanying text and data underlying Table 8.

b That is, excluding equity in or contribution to group pension plans and social security taxes.

c Does not exactly equal difference between figures shown in preceding two columns because of rounding.

tically a highly significant increase.[18] This is a striking confirmation, using separate data and comparing the same group with itself, of the earlier results from this sample. The size of the difference is even larger than found previously, though this may in part reflect bias.

One reason for bias is that the regression ignores any changes in the saving ratio with age. The estimated increase in saving with coverage may have an upward bias if the true saving ratio rises with age, as is likely, and if age and length of time covered are positively correlated. To that extent the groups covered longer have a higher estimated rate of saving, partly because they are older and not only because they are covered. Our estimate of S_c relative to S_{nc} will be too high.[19]

18 If errors in the measurement of wealth are positively skewed, residual errors from the regression may not be normally distributed, in which case the estimates of error are inaccurate.

19 If $P = P_c + P_{nc}$, $P + 22$ is the age of the head of the household. The regression equation in the text may then be written:

$$\frac{W}{Y} = \left(\frac{W}{Y}\right)_{22} + S_c P_c + S_{nc}(P - P_c)$$
$$= \left(\frac{W}{Y}\right)_{22} + (S_c - S_{nc})P_c + S_{nc}P,$$

so that S_{nc} is the average saving ratio over working years and $(S_c - S_{nc})$ is the extra

Another bias of uncertain direction results from using the latest year's income to deflate wealth. If income is constant over time, there is of course no bias. But if income rises over time, as is to be expected, our estimate of the average past saving ratio is distorted in two ways. On the one hand, the average ratio is understated simply because the latest year's income overstates earlier income levels, and this bias increases as the average covers a longer period of time. On the other hand, the later ratios in the average are understated less than the earlier ratios in using the latest year's income. If the true saving ratio rises with age, this progressively smaller understatement makes our estimate of the average past ratio rise even more with age than the true ratio does.[20] If this second distortion exceeds the first, the effect is to

saving when covered. Positive correlation between the two "independent" variables increases the standard error of estimates, but ordinarily does not bias them.

If the true S_{nc} rises with age, however, the assumption of its constancy in the regression equation transfers part of the rise into a higher estimate of $(S_c - S_{nc})$, because P_c tends to be high when P is high. The correlation coefficient between P_c and P for Table 9 is .51—not negligible, but not particularly high either.

[20] Let the true saving-income ratio of a household for the ith year be $\frac{S_i}{Y_i}$. Then the average past saving ratio is

$$\frac{\sum_{1}^{n} \frac{S_i}{Y_i}}{n}. \tag{1}$$

where n is the number of working years of the household. Our estimate of this average is the ratio of wealth to latest year's income divided by n. If all wealth represents saving during working years, the estimate of present wealth to income is

$$\frac{\sum_{1}^{n} \frac{S_i}{Y_n}}{n}. \tag{2}$$

It can be shown that the difference between the two is changed as follows by a one-year increase in n:

$$\underset{n}{\Delta}([1]-[2]) = -\left(1 - \frac{n-1}{n}\,\frac{Y_{n-1}}{Y_n}\right)\frac{\sum_{1}^{n-1} \frac{S_i}{Y_{n-1}}}{n-1} + \frac{\sum_{1}^{n-1} \frac{S_i}{Y_i}}{n-1}\left(\frac{1}{n}\right). \tag{3}$$

If Y is constant, this expression is zero. For a monotonically increasing Y, the expression is positive if $\frac{S_i}{Y_i}$ is the same for all i. Otherwise the expression may be positive or negative, even if both $\frac{S_i}{Y_i}$ and Y_i are rising.

reinforce that of a rise in the true ratio pointed out in the preceding paragraph: our estimate of S_c relative to S_{nc} will be too high to the extent that age and length of time covered are positively correlated.

An offsetting bias is due to our exclusion of extreme wealth-income ratios. Since these would be clustered in the older age groups, their exclusion makes the estimated rate of accumulation of wealth lower, though this is proper if most extreme ratios involve reporting errors or otherwise distort the results.

One other possible bias should be mentioned, though its effect cannot be determined. The regression implicitly assumes that households with the same length of coverage but different ages at the time of the survey had more or less the same saving ratios at the same age. Random variations can be tolerated but not systematic variations. The latter seem likely to be unimportant, but one deserves mention. Since households in different age groups had the same age in different years, the similarity of periods is important. The 65-and-over group was 45–49 in 1938–42, and might have saved more or less then than the present 45–49 group did in the survey year because of differences in the economic climate of the two periods.

Whatever the net effect of these various biases, they could not account for the entire difference between the saving ratios, because Table 8 shows unmistakably that the wealth ratio rises (and hence also the average saving ratio) with length of coverage *within each age group*. (The correlation between age and length of coverage within these narrow age intervals is bound to be negligible.) We may demonstrate the rise in the wealth ratio with length of coverage in Table 8 by estimating the extra saving in covered years within each age group. The estimate is 2.1 per cent per year for all groups combined,[21] the same as estimated in Table 9. Apparently the biases previously discussed cancel each other.

The estimate of saving for the group not covered at the time of the survey is about the same as that for the covered group during the years

[21] This is the least-squares estimate of S_c in the following regression equation:

$$\frac{W}{Y} = S_c P_c + a_1 A_1 + \cdots + a_n A_n,$$

where the a's are regression coefficients and the A's are dummy variables for each group, A_i being unity if the household is in the ith age group and zero otherwise.

it was not covered.[22] We noted earlier, however, that the not-covered households have slightly higher wealth ratios over all than the more recently covered households, which shows up in the regressions as a larger constant term. This might explain why they save less, since greater wealth may dampen the desire to save. We can check the importance of this effect by comparing the saving ratios of the covered and not-covered groups in the same wealth class. Table 10, which makes

TABLE 10

Average Ratios of Total Other [a] Saving to Income of Covered and Not-Covered Households [b] by Total Other [c] Wealth

Total Other Wealth	Average Saving Ratio (per cent) Covered	Not Covered	Difference of Covered over Not-Covered Group[d]
Less than $5,000	5.8	5.7	0.1 (±1.2)
5,000 - 9,999	8.1	7.3	0.8 (±1.4)
10,000 - 19,999	9.2	8.7	0.4 (±1.2)
20,000 - 29,999	10.5	9.4	1.1 (±1.7)
30,000 - 39,999	11.5	10.4	1.1 (±2.3)
40,000 - 59,999	12.6	10.8	1.8 (±2.7)
60,000 and over	14.9	12.7	2.2 (±2.9)
Sum of differences for all groups above			7.6[s] (±5.2)

Note: Figures in parentheses give range of error at the .05 level of significance; s means significantly different from zero at that level.

[a] That is, excluding contributions to group pension plans and social security taxes.

[b] All age classes; same exclusions as for Table 6 (see note c to that table) as well as households for which wealth could not be computed and those reporting negative total other wealth (more debts than assets). Number of households: 10,402.

[c] That is, excluding equity in group pensions.

[d] May not exactly equal the difference between the ratios shown because of rounding.

this comparison, shows that covered households still save more in each class, and that the differences for all classes together are significant. A comparable tabulation, using the ratio of wealth to income and excluding the younger age groups, gives similar results (see Appendix Table H).

[22] Both estimates are too high if the amount of wealth inherited per year rises with age, though as a general rule this seems doubtful.

We should interpret this table cautiously. Since the saving ratio appears to rise with wealth, it is tempting to conclude that wealth has a positive effect on saving. No such conclusion is warranted, however, because wealth depends on past saving, and greater wealth reflects in part a higher propensity to save. To analyze the effect of wealth on saving, one would have to classify by inherited wealth only, or use some other classification that was independent of the propensity to save. Ordinarily, the wealth classification is a way of bringing together households with the same propensity to save as indicated by past saving, particularly if the ratio of wealth to income is used. The problem here is that wealth is also affected by differences in saving between covered and not-covered households: our previous results imply that the households covered for some time ought to have larger wealth and higher wealth ratios. One way to avoid this difficulty would be to compare the saving of households in their first year of pension coverage with not-covered households, both groups stratified by their past saving ratios as indicated by their previous year's wealth-income ratios. The sample unfortunately is not large enough to make such a selective comparison feasible.

How then should we interpret Table 10, in which covered households appear to save more compared with not-covered households that have accumulated the same amount of wealth? If the not-covered households save less currently but have accumulated the same amount of wealth, they may have received greater inheritances or have experienced a slower growth of income in preceding years (see footnote 13), or the greater saving of most covered households may not have lasted long enough yet to make their wealth perceptibly larger. After all, we find that their annual nonpension saving is greater by about 1 per cent of income. Even in ten years this amounts to less than 10 per cent of income, an insignificant (and likely imperceptible) sum for households with a wealth-income ratio over 1.5 or 2.

In summary, the higher other saving of the covered households is associated with their coverage under a pension plan and apparently not with any other characteristic that might cause a higher propensity to save. This conclusion rests on two characteristics of the sample: within age groups the wealth ratio rises with length of coverage, and (as is shown later) the saving ratio does not rise with length of coverage. These two findings imply that the saving ratio steps permanently to a higher level when households become covered.

5. Purchases of Consumer Durables

If pension coverage induces a slight increase in other saving, how do households' budgets accommodate this shift in the disposition of income? Two ways are of interest: the extra saving forces reductions in either durable goods purchases or in current expenses. It would reduce both if the deduction of employees' contributions from paychecks were treated like a reduction in wages and accommodated indiscriminately by an across-the-board cut in all expenditures. But the main impact might fall on current expenses if funds for durable goods were set aside first and not disturbed except for large changes in income. Or, alternatively, the stress might be put on durables if their purchase depended on the residual funds left over after first taking care of current expenses or if, perhaps, durables were considered substitutes for financial assets (both being alternative forms of accumulating "capital"). We may determine the main impact in this sample by comparing the ratio of other saving, including durable goods purchases, to income for covered and not-covered households.

Durable goods purchases over the year were estimated from questions 6, 7, 11, and 12 (see Appendix). The net cash expenditure on cars purchased in the second six months was written in, but on cars purchased in the first six months it had to be estimated from gross price. Each other durable purchase was valued at $300 (arbitrarily selected as the average purchase price). The total of all durable purchases checked was adjusted for seeming duplication (except for cars, furniture, carpets, and hi-fi sets) in cases where the same household checked the same item on both questionnaires. Except for duplication, it seems more accurate to estimate these expenditures from two questionnaires each covering six months than one covering a year, since errors of recall will multiply as the period covered lengthens. It would have been desirable, also, to deduct from the value of purchases the depreciation of durable goods owned during the year, but this was not done. Because of this omission and the nature of the data, these estimates of durable goods purchases are extremely rough and have a large random variability, though their inclusion in the saving ratios offsets variability otherwise present that arises from counting increases in consumer instalment debt as negative saving. The net observed standard error of the ratios is increased by including these purchases, but not by very much. In any event, there is no reason for their inclusion to bias a comparison of covered and not-covered households.

Table 11 presents this comparison of the saving ratio, including and excluding consumer durables, for a subgroup of the sample. It is the

TABLE 11

Average Saving Ratios,[a] Including and Excluding Purchases of Consumer Durables, for Covered and Not-Covered Households[b]

Purchases of Consumer Durables	Average Saving Ratio (per cent)		Difference of Covered over Not-Covered Group[c]
	Covered	Not Covered	
Included	19.0	18.9	0.1 (±0.7)
Excluded[d]	8.7	7.7	1.1[s] (±0.7)

Note: Figures in parentheses give range of error at the .05 level of significance; s means significantly different from zero at that level.

[a] Excluding contributions to group pension plans and social security taxes.

[b] For both rows of table, reduced sample of 7,541 covered households and 2,708 not covered, after excluding self- and not-employed, incomplete questionnaires on durables purchases, and households with unusual gains or losses over $1,000 or with saving (including purchases of consumer durables) greater than 64 per cent or less than −49 per cent of income.

[c] May not exactly equal difference between the ratios shown because of rounding.

[d] Differs from Table 4, bottom two rows, only because of exclusion of questionnaires with no information on durable purchases and of inclusion of some households previously excluded for extreme saving ratios.

same reduced sample used for the preceding tables except that the upper cutoff point for extreme saving ratios (including durables) was put at 64 per cent instead of 49 as before, in order to allow for an average rise in the ratios of fifteen points because of including purchases of durables. (Actually this was an overgenerous allowance, since the average ratio for the full sample rose only about ten points.) The lower cutoff point of −49 per cent was not changed. After this increase of the upper cutoff point, the size of the sample was about the same as before; the necessity of excluding questionnaires with incomplete answers to the questions on purchases, however, reduced the size by about a quarter. This reduction leaves the average difference in saving ratios (excluding durables) between covered and not-covered households nearly the same as in Table 4.

A comparison of the ratios in the table shows that the increase previously noted in discretionary saving from pension coverage comes mostly out of expenditures on consumer durables rather than current expenses. The average saving ratio, including durables, is virtually the same for covered and not-covered households, but if durables are excluded it is slightly and significantly higher for the former. Appendix Table I, which standardizes the ratios for age and income and is comparable to Table 6, except for the inclusion in saving of durables purchases, also shows virtually no effect of coverage on saving.

How then do covered households' budgets accommodate the increase in other financial saving and their own pension contributions? The increase in other financial saving does not show up when combined with durables purchases; hence the increase occurs at the expense of durables expenditures. The cut in the durables budget, to be sure, is relatively small: the average increase in discretionary saving is less than 1 per cent of income and purchases of durables are about 10 per cent, so that on the average they are cut by less than a tenth. Pension contributions out of take-home pay, however, do not decrease the saving plus durables purchases of covered households; hence these contributions occur at the expense of nondurables expenditures—current expenses. It is as though funds to buy durables were budgeted and set aside first, and current expenses were accommodated to whatever then remained of disposable income less the deduction for the pension contribution. There is a clear difference, therefore, in accommodating household budgets to pension contributions, on the one hand, and to the increase in other financial saving, on the other, for what reason is not apparent.

This outline of household budgeting clears away an objection that might be made to our findings. Covered households, being less mobile, are more likely to be homeowners and to have larger families, as indeed is borne out by our sample (see Appendix Table J). Homes and children usually lead to less discretionary saving via cash and securities and to more saving via mortgage payments, insurance, and household durables. Yet we observe almost just the opposite: covered households in our sample have more discretionary saving and purchase fewer durables, though they do increase other contractual saving minutely (see Table 4). Differences between the covered and not-covered households in size and homeownership appear to be too small to eliminate the increase in saving we have attributed to pension coverage or to explain it by differences in income taxes.[23]

One might interpret these results quite differently as showing that the observed increase in discretionary saving is an illusion produced by the method of measuring saving. Saving was initially found by the net change in financial assets, and any debt incurred by a household in acquiring a durable good was deducted from its saving. In a real sense, however, its wealth and saving are not directly affected by this debt, because if we count consumer instalment debt as a liability, we should also count the durable good acquired as an asset. By including purchases of durables in saving, we correct such misrepresentations. If not-covered households purchased the same amount of durables as covered households, but for some reason bought more on credit (though otherwise had the same amount of financial saving), their financial saving ratio would be lower and would account for our results. Such misrepresentations cannot be important in this sample, however, because increases in debt were small on the average and account at most for a small part of the difference in discretionary saving.

By the same token, a similar argument about reductions in bank accounts and government bonds made to purchase durables cannot be dismissed as unimportant. These accounted for nearly all the difference in discretionary saving (see footnote 4, above). Many households undoubtedly save up bank deposits in anticipation of large expenditures on consumer durables, and in the year of purchase their discretionary saving will be abnormally low or even negative. By including such pur-

[23] Appendix Table J shows that covered households within most age groups have 0.1 more dependent children on the average and own 6 per cent more homes. The average covered household, therefore, has an extra income-tax exemption for dependents of (.1 × \$600) \$60 and for mortgage interest of at most \$60 (a \$20,000 mortgage has interest charges of \$1,000 or so per year, and there are 6 per cent extra homeowners among covered households). The extra exemptions total \$120, which reduces the tax by \$30 in the 25 per cent tax bracket. If the saving ratio is 15 per cent, covered households save \$5 per year more on the average because of lower income taxes, far less than we observe in Appendix Tables B and D.

chases in saving we may obtain a truer picture of their normal saving habits. By this reasoning Table 11 shows the true effect of pension coverage—namely, zero substitution.

But this inclusion is not necessarily appropriate for all households or for this study. One household might buy various consumer durables and another rent them or substitute other services, and the two would have quite different saving ratios, including durables; yet both might undertake the same amount of financial saving. Investment in consumer durables does not, after all, provide funds for retirement or for businesses to borrow. Such conceptual problems obscure comparisons, because no hard-and-fast definitions can be drawn. Yet it seems clear that financial assets are the closest substitute for pension funds. While the exclusion of durables will admittedly introduce some inaccuracies into our measure of financial saving, especially if it is meant to represent "long-run accumulation of wealth," because of the large amount and relative infrequency of such purchases and the prevalence of borrowing to make them (negative saving), the degree of error is not likely to be great in a large sample. In any one year some households will be saving up (or paying off previous borrowings: positive saving) for purchases and some will be making them; average saving for the group should be little affected. In a particular year the average could be pushed one way or the other, of course, but it is hard to believe that such purchases could occur so much more frequently in the not-covered group as to account for its statistically significant differences from the covered group. Consequently, the first interpretation of the results—as showing that pension contributions deducted from paychecks substitute for current expenses and that the induced rise in discretionary saving largely comes out of expenditures on consumer durables—seems much the more plausible one.

6. Summary

We find that saving other than through group pensions is not lower for covered households as a whole, which might imply that a pension fund is considered worthless as a substitute for other assets, possibly because of its illiquidity and, if there is no vesting, its uncertain value as a source of retirement income. This interpretation is contradicted, however, by the slightly higher saving in other forms that coverage seems to induce. The additional saving appears small and not in all tests significant at the .05 level, but the evidence, taken as a whole, suggests that it is not a statistical illusion. It shows up in three separate sets of data: reported pension saving, nonpension saving of covered and not-covered households, and past saving as indicated by present wealth. The analysis of the last set suggests that the results cannot be ascribed to a higher propensity to save of those households that acquire coverage. These three sets of data are independent in the sense that they are based on answers to different questions, though the respondents were the same. The possibility of the same bias in the results from all three sets owing to erroneous or mistaken answers is thereby greatly reduced.

Other possible explanations can also be rejected. Differences between covered and not-covered households in number of dependents and homeownership are too small to matter. The large amount of cash and securities saving in the sample, and its slight increase rather than decrease in response to coverage, argues against the possibility that households view pensions as a substitute for other saving and want to make substitutions but cannot do so because other saving is already at an irreducible minimum. Households may still view the pension fund as illiquid, but will hardly see it as worthless.

A suggested reason why premium payments for life insurance seem to behave in a similar way is the following: "Life insurance is likely to be a form of saving especially appealing to families who happen to have a low level of assets relative to income and age as a result of past events. Thus, high insurance is likely to be associated with abnormally high saving relative to resources." [24] Applied to pension coverage, this explanation argues that some households may, because of severe capital losses, become high savers and in consequence seek jobs offering pension plans. For some households, therefore, acquiring coverage is the

[24] F. Modigliani and A. Ando, "The 'Permanent Income' and the 'Life Cycle' Hypotheses of Saving Behavior: Comparison and Tests," in *Consumption and Saving*, II, 161.

Some tabulations of the Consumers Union sample not shown indicate such a positive relation between premium payments and total saving. To me a more appealing explanation is simply that life insurance is one item that most high-saving households want, so that households segregated by the amount of their insurance saving are not randomly distributed by their propensity to save.

consequence of an increase in the propensity to save. Conceivably, such an explanation has some validity for insurance saving, but it seems unsatisfactory for pension saving. For one thing, a comparison of covered and not-covered households with the same wealth in Table 10 was consistent with our other findings and possibly, though not necessarily, inconsistent with the preceding explanation. The main objection is simply that group pension coverage cannot be acquired easily like life insurance. One may have to change jobs, perhaps one's vocation, to acquire a plan if his present employer does not already offer one, and it seems highly doubtful that very many people would want group coverage enough to make such major career changes.

In view of the results in Section 4, above, comparing households covered different lengths of time, the foregoing explanation can be made consistent with our results only by assuming that some event raises a person's propensity to save and at the same time causes him to acquire coverage. It seems more likely that coverage is the initiating event. The accompanying increase in his other saving may then be explained in terms of the attitude of people during their working years toward preparing for retirement. Pension coverage draws attention to the problems of providing for retirement and goes a long way in helping to solve them. It facilitates the rapidly spreading shift to financial means of providing for retirement from the older reliance on family, rental property, and the small family farm or business. But by itself it is apparently found inadequate; the average household supplements it by additional accumulations, mostly in bank accounts and government bonds, at the expense of consumer durables purchases. I shall call this a "recognition" effect of coverage for want of a better term. This effect adds to the appearance of security-mindedness that is said to characterize the modern age, but, since security-mindedness presumably affects covered and not-covered households alike, it cannot be a cause of the recognition effect. I have purposely avoided the term learning effect, as misleading, since pension plans do not appear to teach households much that they did not already know or could not easily find out.

From the various adjustments of the data we may tentatively infer that the recognition effect is stronger among wage earners than salaried employees and does not increase with age or advanced education. The population as a whole, with its larger fraction of wage earners than the Consumers Union sample, ought, based on these results, to reflect the effects of pension coverage even more strongly. This effect, moreover, is not likely to disappear for some time, if ever, because the older and more educated groups that might be expected to apprehend the problems of old age fully even without the aid of a pension plan, and so to presage the behavior of the whole population in the future when the costs and benefits of pension plans are household commonplaces, respond to coverage in the same way and apparently to the same extent as the rest of the sample.

It does not necessarily follow that the recognition effect is the same for all households. Although it seems unaffected by age and education, it could depend on the length of time a household has been covered by a pension plan (which the tabulations in Section 4 did not test) and various other characteristics of the plan. These are questions explored in the next chapter along with other possible differences in behavior among covered households. The findings of the next chapter are consistent with and extend those just reported.